

Stochastic Contracts for Runtime Checking of Component-based Real-time Systems

Chandrakana Nandi
Department of Computer
Science, ETH Zurich
chandrakana.nandi@inf.ethz.ch

Aurelien Monot
ABB Corporate Research
aurelien.monot@ch.abb.com

Manuel Oriol
ABB Corporate Research
manuel.oriol@ch.abb.com

ABSTRACT

This paper introduces a new technique for dynamic verification of component-based real-time systems based on statistical inference. Verifying such systems requires checking two types of properties: functional and real-time. For functional properties, a standard approach for ensuring correctness is Design by Contract: annotating programs with executable pre- and postconditions. We extend contracts for specifying real-time properties.

In the industry, components are often bought from vendors and meant to be used off-the-shelf which makes it very difficult to determine their execution times and express related properties. We present a solution to this problem by using statistical inference for estimating the properties. The contract framework allows application developers to express contracts like “the execution time of component X lies within γ standard deviations from the mean execution time”. Experiments based on industrial case studies show that this framework can be smoothly integrated into existing control applications, thereby increasing their reliability while having an acceptable execution time overhead (less than 10%).

Categories and Subject Descriptors

F.3.1 [Logics and Meanings of Programs]: Specifying and Verifying and Reasoning about Programs; G.3 [Probability and Statistics]: Distribution functions

Keywords

Design by Contract, runtime verification, real-time systems, statistical inference, component-based engineering

1. INTRODUCTION

A real-time system is one for which respecting deadlines is a part of the system specification. These systems are widely used in developing safety critical applications [1] such as control systems for nuclear power plants, safety systems of automobiles, and avionics. Verifying such systems ensures

that their behaviors satisfy the specifications. Among others, component-based engineering is an approach for developing real-time safety critical systems [15], [3], [23] and a lot of research has been done in verifying these systems [16], [8], [17], [9], [12]. Most of these works are based on model checking. Another approach which is better suited for verifying component-based systems due to its support for modular abstraction is *Design by Contract* [18] (annotating programs with executable pre- and postconditions) for specifying functional and real-time properties [5], [22], [4], [14] which can be verified using static or dynamic techniques.

A major challenge in specifying contracts for component-based real-time systems lies in expressing the real-time properties. The existing work in this area focuses mainly on checking very simple timed properties which may not be sufficient for representing the behavior of a system. In practice, real-time contracts for component-based systems should be able to account for all real-time aspects such as availability of resources like CPU and memory, exception handling during the execution of a component and interactions with other components. Another problem lies in reasoning about the Worst Case Execution Times (WCETs) of the components. Static analysis is a possible approach but it is not easy to design and implement a static analyzer that can cater to the needs of a particular industrial component-based system. Even if one does so, it would not be reusable for other such systems. Moreover, static analysis may not be feasible when the components are bought from vendors and the source code is not available [21]. Another approach is to use probabilistic analysis for determining the probability density function of the execution times. This is not applicable to all systems because the execution time distribution depends on the behavior of a particular component. The results for one component can not be generalized to all components.

In this paper we present solutions to the problems of expressing and checking real-time properties by incorporating empirical stochastic analysis in the contracts: instead of determining the exact WCET values of the components or predicting the execution time probability distributions, we *estimate* the WCET values based on empirical distribution functions and then specify the contracts. Our contract framework allows us to specify properties like “**Execution time of component X should lie within γ standard deviations from the mean execution time of X** ” and check them dynamically. The mean and the standard deviation of the execution times are estimated by statistical inference and the parameter γ is a property of a particular component.

We validated our approach on a state-of-the-art component-based framework, FASA (Future Automation System Architecture) [20] developed at ABB Corporate Research. Our experiments are based on industrial case studies and show that for applications deployed on a single host controller, the average execution time overhead added due to the contracts is less than 10%, which makes it efficient and easy to incorporate on top of any component-based real-time application.

2. OVERVIEW

2.1 Motivating Example

We consider a typical control loop with a combination of a feed-forward and a feedback controller as our running example. Figure 1 shows the schematic diagram for this system. The application has a component, Sensor which measures




Figure 1: A control system with a combination of an open loop and a closed loop controller

the value of a physical quantity, in our case electric current generated by System. The feedback loop in our example is based on negative feedback because the output current measured by Sensor is subtracted from the desired current value to generate an error. This error is used by the component, Feedback controller to dynamically regulate the output of the System. The other component is a Feed-forward controller. This controller does not measure the output current generated by System and as a result, the latter has no effect on it. The net effect on the input to System is the sum of the effects of the two controllers. In practice, a combination of a feed-forward and a feedback controller is often used for ensuring faster convergence to the desired value of the physical quantity being measured.

2.2 FASA

FASA (Future Automation System Architecture) is a component based framework developed at ABB Corporate Research [20] used for *cyclic* real time control applications. A FASA application is made up of *function blocks* (components) which communicate with each other through *ports*. Ports can be input or output. A function block receives data at an input port and sends out data through an output port. An input port is connected to an output port through a unidirectional *channel*. FASA has a centralized *scheduler* which is responsible for scheduling the function blocks on the available controllers. The execution of an application is triggered by the launching of the FASA *kernel* which is the main entry point to the framework. The function blocks are implemented as C++ classes and have a dedicated method, `operator()` where their behavior is described.

Referring to our running example, the Feed-forward controller, the Feedback controller, the Sensor, the two Addition Components (shown in Figure 1 with + sign) and the

System are FASA function blocks. The arrows indicate the direction of data flow through the channels.

2.3 Types of Contracts

Our contract framework facilitates the specification and monitoring of functional and real-time properties. For the functional specifications, it supports preconditions, postconditions and class invariants. It supports loop invariants and loop variants for checking the correctness of loops and the `old` construct as used in some languages such as Eiffel [18], which specifies properties about the state of an object. Figure 2 is a code snippet showing a functional contract for Sensor from our running example. Sensor has an attribute `interval` which records the current cycle number. In every cycle, this value is incremented by one. Before every update, the current value of `interval` is saved using the `OLD` macro. In the postcondition, the old value is retrieved using another macro, `GET_OLD` which takes two arguments, the type and the attribute name. The postcondition states that the value must be correctly updated in each cycle.

```
class SENSOR_BLOCK
{
public:
    int interval;
    void operator()
    {
        OLD(interval);
        interval = interval + 1;
        ...
    }
    void functional_post_conditions()
    {
        POSTCONDITION(interval == GET_OLD(int, interval) + 1);
    }
};
```

Figure 2: Functional contract for monitoring the state of a class attribute for Sensor from the running example

For the real-time specifications, the properties supported by our framework are related to:

1. Execution times of the function blocks.
2. Cycle times of the control applications: our target applications are cyclic in nature and since they are safety critical, they must not violate the desired cycle time.
3. Jitter margin: applications must not exceed the maximum deviation from the cycle time that can be tolerated. This upper limit is called *jitter margin*.

The cycle time and jitter margin for industrial component-based systems are usually predefined by the control engineer. The execution time analysis of the individual components is very difficult due to several reasons. First, developing a comprehensive static analyzer for WCET analysis which can be used for any component-based system is very difficult. Second, components are often bought from vendors and used off-the-shelf due to which their source code may not be available. Third, some existing probabilistic approaches try to model the execution times of all components using a particular distribution function; we executed 24 different function blocks from 11 control applications and found that upon executing the blocks for 1000 cycles or more, the distributions converge to a single peak but are very different from one another. As a result, it is not possible to fix a standard probability density function a-priori to model the real-time behavior of all the function blocks. To overcome these

challenges, instead of using static analysis or a particular probability distribution function, we statistically *estimate* the execution time upper bounds and specify contracts using these estimates. An example of a real-time contract for a component in our approach reads like “Execution time of component X should lie within γ standard deviations from the mean execution time of X ”, where the execution time upper bound is computed as a function of the estimates of the mean and the standard deviation. Figure 3 shows this contract for Sensor. Our approach can be used with any component-based system and its performance is independent of the size of the application (in terms of the number of components).

```

struct Real_Time
{
    float estimated_mean;
    float estimated_variance;
    float gamma;
};

class SENSOR_BLOCK
{
public:
    int interval;
    Real_Time real_time;
    void operator()
    {
        UD(interval);
        interval = interval + 1;
        ...
    }
    void real_time_post_conditions()
    {
        POSTCONDITION(execution_time <= real_time.estimated_mean
            + real_time.gamma * sqrt(real_time.estimated_variance));
    }
};

```

Figure 3: Stochastic real-time contract for Sensor from the running example

3. STOCHASTIC CONTRACTS

We use statistical inference for analysing the execution times of the function blocks and specifying related contracts: we estimate the WCET values based on the empirical cumulative distribution functions (cdfs) of the execution times. By definition, the cdf of a random variable X is given by:

$$F(x) = P[X \leq x] \quad (1)$$

where $P[X \leq x]$ is the probability that the random variable X has value less than or equal to x . Let μ and σ represent the mean and standard deviation of the unknown probability density of the execution time of a function block. Using equation (1) we can find an upper bound τ_π on the execution time such that $P[X \leq \tau_\pi] = \pi$ where $\pi \in [0, 1]$. Now, τ_π can be expressed as a function of μ and σ as:

$$\tau_\pi = \mu + \gamma\sigma \quad (2)$$

where γ gives us the distance of τ_π from the mean, μ , in terms of number of standard deviations, σ s. Every function block is executed for n cycles and the execution times are recorded. Let $e_1, e_2, e_3 \dots e_n$ denote this sample. The sample data is used for determining the empirical cdf. Figure 4 shows the empirical cdf of our Sensor, based on 1000 cycles. Next, the mean, μ_n , and the standard deviation, s_n of the sample are computed as:

$$\mu_n = \frac{\sum_{i=1}^n e_i}{n}$$

$$s_n = \sqrt{\frac{\sum_{i=1}^n (e_i - \mu_n)^2}{n - 1}}$$




Figure 4: Empirical cumulative distribution function of the Sensor from the running example based on sample size of 1000

Using equation (1) our framework *estimates* an upper bound, $\hat{\tau}_\pi$ on the execution time, if a threshold probability $0 \leq \pi \leq 1$ is given by the application developer. For example, if the developer sets π to 0.99, then $\hat{\tau}_\pi$ is computed so that 99% of the execution times lies within $\hat{\tau}_\pi$ from the $(n + 1)^{th}$ cycle onward. Thus, the accuracy of the bound depends on π . Similar to equation (2), $\hat{\tau}_\pi$ can be expressed as a function of μ_n and s_n as:

$$\hat{\tau}_\pi = \mu_n + \gamma s_n \quad (3)$$

The reason why we express the bound as a function of μ_n and s_n is that it makes dynamic updating of the bound (as explained later) faster. The framework computes γ only once and treats it as a constant while updating μ_n and s_n and re-estimating the bound. $\hat{\tau}_\pi$, μ_n and s_n are *unbiased estimators* [11] of τ_π , μ and σ respectively. This means that the expected value of $\hat{\tau}_\pi$ is equal to τ_π and similarly for the other two parameters. Following is the formal definition of this notion.

Definition. Let $e_1, e_2, e_3 \dots e_n$ be a random sample drawn from a population and let θ be an unknown parameter of the population which is to be estimated. Let $\hat{\theta} = v(e_1, e_2, e_3 \dots e_n)$ be a function of the sample. By definition, $\hat{\theta}$ is an unbiased estimator of θ if the expected value of $\hat{\theta}$ is equal to θ , i.e.

$$E[\hat{\theta}] = \theta \quad \square$$

We use unbiased estimators because they have desirable properties [11]. The proof of the unbiasedness of μ_n and s_n are standard results in statistical inference. They are also provided in [19]. We use them in the following theorem.

Theorem. For any given function block, $\hat{\tau}_\pi$ is an unbiased estimator of the upper bound τ_π on its execution time.

Proof. $E[\hat{\tau}_\pi] = E[\mu_n + \gamma s_n]$
 $= E[\mu_n] + E[\gamma s_n]$
 $= E[\mu_n] + \gamma E[s_n]$
 $= \mu + \gamma\sigma$
 $= \tau_\pi$ using equation (2) \square

Algorithm 1 describes the process for computing γ . It takes the execution time data and π as inputs and returns γ . The algorithm sorts and bins the execution times, where the number of bins is the square root of the total data count and calculates the upper limit for the first bin (line 12). Then it computes the probability of the data falling into the first bin (lines 13-20). If this probability exceeds (or is equal to) π , $\hat{\tau}_\pi$ is set to the upper limit of the first bin. Otherwise, it computes the next bin's upper limit (line 25)

and the process repeats until π is reached. Ultimately, the algorithm returns γ computed using equation (3).

Algorithm 1 Empirical Cdf

```

1: procedure EMPIRICALCDF(execution_times,  $\pi$ ):  $\gamma$ 
2:    $sum, y \leftarrow 0, 0$ 
3:    $buffer\_size \leftarrow \text{SIZE}(\text{execution\_times})$ 
4:    $\text{SORT}(\text{execution\_times})$ 
5:    $min \leftarrow \text{execution\_times}[0]$ 
6:    $max \leftarrow \text{execution\_times}[buffer\_size-1]$ 
7:    $\mu_n \leftarrow \text{MEAN}(\text{execution\_times})$ 
8:    $s_n \leftarrow \text{STD\_DEV}(\text{execution\_times})$ 
9:    $upper\_bound \leftarrow max$ 
10:   $number\_of\_bins \leftarrow \text{SQRT}(buffer\_size)$ 
11:   $bin\_width \leftarrow \frac{(max-min)}{number\_of\_bins}$ 
12:   $upper\_limit\_of\_bin \leftarrow (\text{execution\_times}[0] + bin\_width)$ 
13:  while  $upper\_limit\_of\_bin \leq max$  do
14:    while  $y \leq buffer\_size$  do
15:      if  $\text{execution\_times}[y] \leq upper\_limit\_of\_bin$  then
16:         $sum \leftarrow sum + 1$ 
17:      end if
18:       $y \leftarrow y + 1$ 
19:    end while
20:     $probability \leftarrow \frac{sum}{buffer\_size}$ 
21:    if  $probability \geq \pi$  then
22:       $\hat{\tau}_\pi \leftarrow upper\_limit\_of\_bin$ 
23:      break
24:    end if
25:     $upper\_limit\_of\_bin \leftarrow upper\_limit\_of\_bin + bin\_width$ 
26:     $sum, y \leftarrow 0, 0$ 
27:  end while
28:   $\gamma \leftarrow \frac{\hat{\tau}_\pi - \mu_n}{s_n}$ 
29:  return  $\gamma$ 
30: end procedure

```

Once γ (based on the data from the first n cycles) is obtained, we can specify contracts like “Execution time of function block X should not exceed γ standard deviation from the mean of X ” from the $(n + 1)^{th}$ cycle onward. The parameter γ is a characteristic property of a particular function block. This technique of computing a function block specific parameter γ and using it for specifying the real-time contracts eliminates the problems of using the same probability density function to model the execution times of all function blocks.

As the function blocks continue to execute beyond the first n cycles, the accuracy of the upper bound may deteriorate due to several reasons. First, a function block may take longer time to execute due to platform related reasons such as waiting to acquire a resource and this might affect its execution time. Second, the execution time depends on the path of execution followed by a function block: in a certain cycle, it may have to handle an exception which might take longer than usual. Third, the execution time of a function block depends on interactions with the other function blocks and if the other function blocks have a delay, then that will affect the execution time of all the interacting function blocks. To account for these factors, it is important to dynamically update the computed upper bound from time to time. Since all computations are performed at runtime, running Algorithm 1 too many times during the execution of the control application would degrade the application’s performance. To prevent this, we only re-estimate μ_n and s_n at runtime and update $\hat{\tau}_\pi$ using the new value of μ_n and s_n and the previously computed value of γ (based on the first n cycles). For updating the estimates, we use a sliding window

mechanism: after a certain number h of cycles, the oldest h values of the execution time of a function block are replaced by the newest h values and μ_n and s_n are recomputed.

4. IMPLEMENTATION

We implemented the contract framework as a C++ library. Contracts are checked dynamically. When a contract fails, a message is logged. The logging is done by the FASA scheduler during the time left in each cycle after all function blocks are executed. This time is called the *slack time*. If there is no slack time remaining in a cycle, the messages are stored in a buffer and carried on to the next cycle, where they are logged before the cycle’s own messages. The contracts are specified within the function blocks in “contract methods”. This can be seen in the examples in Figure 2 and Figure 3. To test the performance of our framework, we fixed an upper bound on the execution time overhead added due to contracts as 10%, consistent with standard practices [13].

5. VALIDATION AND RESULTS

We evaluated our framework on five industrial case studies developed at ABB Corporate Research. The applications were tested on MacMini computers with 4 GB RAM and quad code processors, each core running at 1.8 GHz. The operating system was 64 bit Ubuntu 13.04. The results are shown for 10000 executions of each application. For the real-time contracts, the cycle time was taken as 10ms for the first four case studies and 100ms for the fifth case study in order to handle network communication delays and the jitter margin was set to 0.1ms for all case studies. The entire list of contracts for each application is available in [19].

Simple Counter. There is a single function block. It does not have any ports. An integer variable **counter** is initialized to 1 in the block constructor and is incremented in every cycle up to 10. For functional contracts, the framework checked a loop variant, a loop invariant and a postcondition to ensure the correct incrementation of **counter** up to 10 in every cycle. For the stochastic contracts, the value of the threshold probability π was set to 0.95 and the sliding window updating interval h was set to 1.

Gaussian Generator. The application generates random numbers having a Gaussian distribution. A Random Generator block generates two random numbers according to Uniform Distribution, $U[0, 1]$. It then sends the two values to a Gaussian Generator block which generates two standard normal, $N(0, 1)$ random values using the Box-Muller transformation¹. These two values are sent to a Range Calculator which computes the range of the two Gaussian random numbers. The framework checked functional preconditions for verifying the connectivity of the ports for every function block before they started transmitting data. For the Random Generator, it checked postconditions to ensure that the generated random numbers lie within $[0, 1]$, as required by the Gaussian Generator. For the Range Calculator, a postcondition ensured that the computed range is ≥ 0 . The values of π and h were set to 0.97 and 10 respectively for the stochastic properties.

Energy Pack Core Model. This is our running example, shown in Figure 1. The application is described in section 2.1. The functional contracts included postconditions

¹http://en.wikipedia.org/wiki/Box-Muller_transform

Table 1: Performance summary for five case studies

Case Study	Number of contracts	execution time		overhead
		without contracts	with contracts	
Simple Counter	9	0.5447	0.5574	2.33%
Gaussian Generator	27	1.4202	1.5161	6.75%
Energy Pack Core Model*	36	1.7467	1.8657	6.81%
Binary Search*	43	1.5865	1.6761	5.65%
Net Proxy: sender	10	0.4746	0.6359	33.99%
Net Proxy: receiver	10	0.8016	1.1531	43.85%

*median execution time is computed due to outliers in the data.

for monitoring the states of variables and attributes, preconditions for ensuring the connectivity of ports and class invariants. For this application, π was set to 0.95 and h to 5.

Binary Search. A Random Integer Generator generates a random integer which is sent to an Array Creator. The latter stores the integer in a dynamic array. It sends the array to a Sorter to be sorted in ascending order. The sorted array is sent to a Binary Search block. This block has second input port where it receives a random number from the Random Integer Generator and searches for this random number in the sorted array that it receives from the Sorter and shows the index of the number if it is found, otherwise it prints -1 as the index. After every 10^{th} cycle, the memory allocated to the array is cleared. Thus, the maximum size of the array is 10. The functional contracts checked preconditions for the connectivity of the ports and loop variants and invariants to ensure the total correctness of the loops in the Array Creator, the Sorter and the Binary Search block. For the execution time related contracts, π was set to 0.98 and h to 5.

Net Proxy. There are two separate applications launched on two separate host controllers. The first application has a Sender and a Net Proxy Send block. The latter transmits integer values from the Sender through the network. The second application has a Net Proxy Receive block and a Receiver. The Net Proxy Receive block receives the integer value from the network and sends it to the Receiver. In order to synchronize the clocks on the two hosts, the Precision Time Protocol (PTP) is used. For the Sender, the functional contracts included a precondition to check the connectivity of the output port, a postcondition to monitor the state of an attribute and a class invariant. For the Receiver, a precondition checked the connectivity of the input port and a postcondition monitored the state of an attribute. π was set to 0.99 and h to 10 for both applications. A real-time postcondition checked the completion time of the Receiver with respect to the Sender in a contract like “Execution of Receiver must terminate within 0.2ms from the start of the execution of the Sender”, in order to account for the network communication delays.

Table 1 shows the total number of the contracts, the mean execution times in ms for the five case studies with and without contracts enabled and the execution time overheads. The number of contracts include all functional and real-time contracts for every function block in the respective application [19]. The execution time overhead for the first four applications is well under our 10% limit. For the fifth case study, the overhead was more due to delays in the network communication (some of the real-time contracts for the *Receiver* relied on data obtained through the network).

ceiver relied on data obtained through the network).

6. RELATED WORK

[25] and [6] describe use of contracts for verification of programs in Ada. In [7] a tool has been developed for the BIP [3] component framework which checks a program against programmer written specifications at runtime. [5] introduces a framework for specifying contracts in temporal logic. The use of contracts for distributed embedded system design is shown in [27]. None of the above work supports real-time contracts for reasoning about metric time. Closest to our

Table 2: Characteristics of different contract frameworks

	functional	temporal	stochastic updates
AdaCore[25]	✓	×	×
Hi-Lite[6]	✓	×	×
Barbacci et al.[2]	✓	✓	×
Härtig et al.[10]	✓	✓	×
Stierand et al.[27]	×	✓	×
Sangiovanni-Vincentelli et al.[24]	✓	×	×
RV-BIP[7]	✓	×	×
OCRA[5]	✓	×	×
Sojka et al.[26]	×	×	×
Stochastic Contracts	✓	✓	✓

idea are the works [10] and [2] in the sense that both focus on specification of functional and timed properties for real-time systems, but, unlike their tools, our tool does not rely only on statically defined contracts. Instead, it updates the contracts at runtime to ensure that they remain meaningful throughout the execution of an application, taking into consideration the changes in the underlying platform related factors. [26] and [24] show contracts in multiple layers of real-time systems. The similarity with our work lies in the layered structure of the contracts. In our framework, we specify real-time contracts in two hierarchical layers: function block level (WCET related) and application level (cycle time and jitter related). [24] only supports functional properties while [26] supports properties related to resource reservation.

To our knowledge, our framework is the first to make use of statistical techniques for dynamically estimating and updating real-time contracts related to execution times of an application’s components. Table 2 summarizes the characteristic features of the existing contract frameworks which are relevant for our research and the characteristics of our own contract framework (the final row).

7. CONCLUSIONS AND FUTURE WORK

We presented a statistical inference based approach for computing real-time contracts for component-based real-time control applications. Based on experiments with industrial control applications deployed on a single controller (without network communication), we demonstrated that our contract framework adds an acceptable execution time overhead (much less than 10%).

We used statistical inference for estimating upper bounds on the execution times of the function blocks. A possible extension would be to use machine learning techniques such as neural networks and compare the quality of the obtained estimates from the two approaches. The challenge in the latter approach lies in limiting the overhead added by the heavy computations at runtime, either by performing the learning offline or by other optimization techniques.

8. ACKNOWLEDGEMENTS

We would like to thank Viktor Kuncak for his guidance and Carlo A. Furia, Chris Poskitt and Nadia Polikarpova for their feedback on the drafts of the paper. The first author is currently funded by ERC grant no. 291389.

9. REFERENCES

- [1] R. Alur. Techniques for automatic verification of real-time systems. PhD thesis, Stanford University, 1991.
- [2] M. Barbacci and J. M. Wing. Specifying functional and timing behavior for real-time applications. In *Proc. Parallel Architectures and Languages Europe, Volume I*, pages 124–140, London, UK, UK, 1987. Springer-Verlag.
- [3] A. Basu, L. Mounier, M. Poulhiès, J. Pulou, and J. Sifakis. Using bip for modeling and verification of networked systems – a case study on tinyos-based networks. *IEEE NCA’07*, pages 257–260.
- [4] S. S. Bauer, A. David, R. Hennicker, K. Guldstrand Larsen, A. Legay, U. Nyman, and A. Wasowski. Moving from specifications to contracts in component-based design. *FASE’12*, pages 43–58, Berlin, Heidelberg. Springer-Verlag.
- [5] A. Cimatti, M. Dorigatti, and S. Tonetta. Ocr: A tool for checking the refinement of temporal contracts. *IEEE ASE’13*, pages 702–705.
- [6] Y. M. Cyrille Comar, Johannes Kanig. Integrating formal program verification with tesing. 2014 http://www.adacore.com/uploads_gems/Hi-Lite_ERTS-2012.pdf.
- [7] Y. Falcone, M. Jaber, T.-H. Nguyen, M. Bozga, and S. Bensalem. Runtime verification of component-based systems. volume 7041 of *LNCS*, pages 204–220. Springer Berlin Heidelberg. In *SEFM’11*.
- [8] G. Gössler, S. Graf, M. Majster-Cederbaum, M. Martens, and J. Sifakis. An approach to modelling and verification of component based systems. volume 4362 of *LNCS*, pages 295–308. Springer Berlin Heidelberg. In *SOFSEM 2007*.
- [9] Z. Gu and K. G. Shin. Model-checking of component-based event-driven real-time embedded software. *IEEE ISORC’05*, pages 410–417.
- [10] H. Härtig, S. Zschaler, M. Pohlack, R. Aigner, S. Göbel, C. Pohl, and S. Röttger. Enforceable component-based realtime contracts. *Real-Time Syst.*, 35(1):1–31, Jan. 2007.
- [11] R. Hogg and C. A. McKean, J.W. *Introduction to Mathematical Statistics*. Pearson Education, 2005.
- [12] P.-A. Hsiung, W.-B. See, T.-Y. Lee, J.-M. Fu, and S.-J. Chen. Formal verification of embedded real-time software in component-based application frameworks. *APSEC’01*, pages 71–78.
- [13] X. Huang, J. Seyster, S. Callanan, K. Dixit, R. Grosu, S. A. Smolka, S. D. Stoller, and E. Zadok. Software monitoring with controllable overhead. *Int. J. Softw. Tools Technol. Transf.*, 14(3):327–347, June 2012.
- [14] R. Koymans. Specifying real-time properties with metric temporal logic. *Real-Time Systems*, 2(4):255–299, 1990.
- [15] G. Lipari, P. Gai, M. Trimarchi, G. Guidi, and P. Ancilotti. A hierarchical framework for component-based real-time systems. *Electronic Notes in Theoretical Computer Science*, 116(0):253 – 266, 2005. TACoS 2004.
- [16] G. Macariu and V. Cretu. Timed automata model for component-based real-time systems. *IEEE ECBS 2010*, pages 121–130.
- [17] G. Madl, S. Abdelwahed, and G. Karsai. Automatic verification of component-based real-time corba applications. *IEEE RTSS’04*, pages 231–240.
- [18] B. Meyer. *Object-Oriented Software Construction, 2nd Edition*. Prentice-Hall, 1997.
- [19] C. Nandi. Contracts for Real-Time, Safety Critical Systems, 2014. M.Sc thesis, EPFL, Switzerland. <http://infoscience.epfl.ch/record/203686?ln=en>.
- [20] M. Oriol, M. Wahler, R. Steiger, S. Stoeter, E. Vardar, H. Koziolk, and A. Kumar. Fasa: A scalable software framework for distributed control systems. *ACM ISARCS’12*, pages 51–60, New York, NY, USA.
- [21] R. Perrone, R. Macedo, G. Lima, and V. Lima. An approach for estimating execution time probability distributions of component-based real-time systems. *J. UCS*, 15(11):2142–2165, 2009.
- [22] S. Quinton and S. Graf. Contract-based verification of hierarchical systems of components. *SEFM’08*, pages 377–381.
- [23] S. Richter, M. Wahler, and A. Kumar. A framework for component-based real-time control applications. In *13th Real-Time Linux Workshop, Prague, Czech Republic*, 2011.
- [24] A. Sangiovanni-Vincentelli, W. Damm, and R. Passerone. Taming dr. frankenstein: Contract-based design for cyber-physical systems*. *European Journal of Control*, 18(3):217 – 238, 2012.
- [25] E. Schonberg. Towards Ada 2012: An interim report. *ACM SIGAda ’10*, pages 63–70, New York, NY, USA.
- [26] M. Sojka and Z. Hanzalek. Modular architecture for real-time contract-based framework. *IEEE SIES ’09*, pages 66–69.
- [27] I. Stierand, P. Reinkemeier, T. Gezgin, and P. Bhaduri. Real-time scheduling interfaces and contracts for the design of distributed embedded systems. *SIES’13*, pages 130–139.